

Probe Design for Tracking Based on LED Imaging



Larry D. Davis, Jr
Jannick P. Rolland
Yohan Baillot

TR98-003

Center for Research and Education in Optics and Lasers (CREOL)
Department of Electrical and Computer Engineering
Department of Computer Science,
4000 Central Florida Blvd
Orlando, FL 32816-2700

1. Introduction

There is a distinct need in virtual environment applications for the accurate tracking of user motion. This need is emphasized for Augmented Reality applications because virtual objects must be properly registered with respect to real objects in the environment [1]. Within Augmented Reality (AR), applications themselves can be more or less demanding. Medical applications, for example, can require accuracy and resolution to be on the order of a fraction of a millimeter [2][3].

While system delay, or lag, is a major source of registration error in AR environments [4], a first requirement for accurate registration (and most generally tracking) is accurate and precise static measurements of points and objects in the virtual environment. This paper presents static methods of probe design for optical tracking as a first step towards registration in dynamic environments. Although dynamic registration is not explicitly treated in this paper, frame rate requirements to minimize tracker lag are accounted for in the design specification of the probes. The system lag, which often limits the overall performance, is not treated in this paper.

To optimize tracking accuracy and resolution in our augmented reality research, we employ optical tracking technology because contemporary optical tracking systems can typically provide accuracy and resolution in the sub-millimeter range. Also, optical tracking provides fast, accurate position and orientation data. In addition, optical tracking is not subject to electromagnetic interference or acoustic noise, an important requirement in most applications and especially important for the high performance systems required in augmented reality.

A common approach to optical measurement of position and orientation relies upon an arrangement of beacons in the environment [5][6]. These beacons are most commonly activated sequentially, but could also be activated simultaneously via technology choices. In either case, they are observed by a camera or some other optical sensing device. In some configurations, the positions of the beacons are fixed and the camera or set of cameras moves[7]. Other configurations have stationary cameras with beacons arranged on a probe and attached to a mobile target. The work presented here focuses upon the latter configuration, where probes are positioned on a moving object whose position and orientation must be determined.

We present a framework for designing probes based on tracking requirements. We first review the working principles of the tracking system configuration selected and provide specifications of the Northern Digital OPTOTRAK 3020 tracking system used in the implementation of the work. We then provide general guidelines for selecting a preferred probe geometry and a method for the design of spherical probes. Next, we provide a methodology to predict accuracy and resolution in both position and orientation of a probe given its geometry, number of beacons, and the accuracy and resolution of a single beacon. While our specific aim is to design high accuracy and resolution probes for augmented reality with application to medical visualization [8][9], the methodology proposed is generally applicable. Finally, we demonstrate the basic principles by applying the framework to four probes sensed by the OPTOTRAK.

2. Fixed Camera Tracking Configurations

First, we define a probe as a rigid configuration of LEDs which is attached to a real object to determine position and orientation. A fixed camera configuration includes a collection of beacons usually assembled in a probe or set of probes, a controller to activate the beacons (if necessary), sensing devices, and processing units to provide position and orientation data of probes from the individual beacon emissions. A probe is placed on a target object and each beacon emits light energy the sensing devices are attuned to. Beacons are typically differentiated from each other based upon either their activation sequence or their frequency of activation. If two or more sensing devices detect at least three beacons on a probe, the orientation and position of the probe, and therefore the target object, can be determined.

2.1 Implementation: OPTOTRAK 3020

The OPTOTRAK 3020, an off-the-shelf, fixed camera optical tracking system, is used in validating the proposed probe design framework. Its operation is based upon detecting the positions of sequentially activated, infrared light-emitting diodes in plastic housings, known as beacons. When each beacon is activated, three vertically aligned cameras in the OPTOTRAK position sensor record its position, provided the beacons are within the working volume (26.98 m³) and are visible to the cameras. Each position is compared to a predefined spatial arrangement of beacons, called a rigid body. If the cameras are able to detect three of the beacons on the rigid body, the position and orientation of the probe may be determined [10].

Similar to other optical tracking devices, the OPTOTRAK has a high update rate without environment interference. It can provide real-time 3D data at frame rates up to 600 Hz [10]. Also, the OPTOTRAK system is flexible, in that the main body housing the cameras may be operated in multiple orientations; this allows some relief from the potential occurrence of occlusion between the beacons and the cameras. Finally, the system is calibrated in-factory and is robust enough that additional calibration is not usually required.

For all its advantages, however, having the OPTOTRAK only insures the possibility of state-of-the-art tracking. The technology was not developed for use in virtual environments and, therefore, specific probes need to be designed for head tracking and the tracking of human motion in general. The framework presented here allows optimization of probe design for specific tracking requirements given application constraints.

3. Design Approach

In designing a probe, there are several key factors which will impact the results. These factors include the field of view (FOV) of tracker and the beacons, the working volume of the tracker, the permissible frame rate, and the types of user interaction in the virtual environment. For a given probe we define its geometry to be how the beacons are placed, the distance between the beacons, and the number of beacons used

First, methods for determining beacon configuration as a function of the types of user interaction in the virtual environment are discussed. Included in this discussion is the tradeoff between the

number of beacons used and the achievable frame rate. Then, methods of designing a probe given the tracker FOV are discussed.

3.1 Beacon Configurations

Probes may differ in their intrinsic geometries. For a given probe geometry, the optimal number of beacons used must be determined based upon frame rate requirements, cost, and the number of beacons required by the tracker. Similarly, a given number of beacons will suggest possible geometries. The simplest probe design involves one beacon. However, only position data may be obtained with one beacon.

A planar configuration is the simplest geometry that can provide position and orientation data. While three beacons are sufficient to define a plane, additional beacons may be necessary to meet probe resolution requirements, as demonstrated in section 4. Another consideration for increasing the number of beacons on a planar probe (in rare cases) is to ensure the sensing devices always detect three beacons at any time. For some applications where user motion through the environment is significantly limited, users can be approximately oriented toward the cameras, and a planar assembly may provide adequate performance.

Applications involving large angular displacements within the working volume require non-planar probe geometries. To facilitate the increased range of motion additional beacons must be provided on non-planar probes. Placing additional beacons on the probe, however, will lower the effective frame rate of the tracker. In a sequential scheme, the relationship between the desired frame rate and the number of beacons is given by

$$FrameRate = \frac{f_a}{N}, \quad (1)$$

where f_a is the frequency at which the beacons are activated and N is the number of beacons.

This design tradeoff between greater range of motion and desired frame rate constrains our solution to one that offers a high degree of symmetry. By determining a symmetric distribution of beacons about the centroid of the probe, we can insure tracking over a large field of regard with a minimum number of beacons. A spherical probe geometry will satisfy these requirements. However, based on the number of beacons required, other geometries may be considered to shorten the design process considerably.

Non-spherical, symmetric geometries point to five spherical equivalents, known as platonic solids. These solids are the tetrahedron, octahedron, cube, icosahedron, and dodecahedron, which have 4, 6, 8, 13, and 20 vertices, respectively. These polygons are spherical equivalents because for each representation, a sphere can be specified such that all the vertices of the solid satisfy its equation. By placing beacons at the vertices of a platonic solid with the correct radial distance, full orientation tracking is possible with the appropriate placement of a number of sensing devices in the environment. In that case, the number of vertices on the solid imposes the number of beacons. Therefore, given a number of beacons, the platonic solid with a number of vertices

closest to the number of beacons may be selected for optimal symmetry with no further need for design.

If a platonic solid cannot be used to meet tracking requirements, a spherical geometry is required. In this case, the beacons must be distributed equally on a sphere to allow the largest field of regard with the fewest number of LEDs. To determine how to uniformly distribute the beacons on the sphere, we can model beacons on the sphere as atoms in a molecule. Continuing in this manner, an algorithm for uniform distribution can be based on minimizing the energy between the atoms, similar to the way a molecule naturally reaches a minimum energy state. Therefore, we treat each beacon on the probe as a charged particle and disperse the particles on a spherical surface using electromagnetic force laws. To find the minimum energy state, we use the Metropolis variation of the simulated annealing optimization approach [11]. A flowchart detailing this approach and results obtained with the algorithm can be found at the end of the document. The construction of a spherical probe using this algorithm is currently in progress and will be discussed at a later date.

3.2 Relationship of Probe Design to the Tracker FOV

Trackers of the configuration discussed here have a limited working volume, determined in part by the field of view (FOV) of the sensing devices, in which tracking can be performed. Moreover, the beacons have a conic emission pattern (referred to as the beacon solid angle) that must fall within the tracker FOV to be detected. Therefore, the spacing of the beacons will be constrained to meet the requirement that three beacons must be detected simultaneously by the sensing devices. This requirement will add a constraint to the possible scale of the probe to allow continuous tracking in orientation. However, this constraint may yield a decrease in resolution in orientation as described in section 4.4. Therefore, beacon spacing yields a tradeoff between continuous tracking in orientation and achievable resolution.

Another solution to limited tracker FOV is to increase the number of beacons on the probe. Depending upon the application, though, the loss in achievable frame rate, as described in the preceding section, may be unacceptable.

4. Modeling the Accuracy and Resolution of Probes

Given the position of a probe in 3-D space, accuracy in position or orientation is defined as the error between the known and measured position values averaged over a large number of tracker samples. Resolution is defined as the uncertainty in position or orientation of the probe between tracker samples. The accuracy and resolution of a probe, regardless of its geometry, depends upon the accuracy and resolution of a single beacon. We present a static method for predicting accuracy and resolution of a probe as a function of its number of beacons and the resolution and accuracy of a single beacon

The importance of a static model is that it leads to a dynamic model. By measuring the accuracy of the probe statically, an estimate of the bias induced by the tracker can be computed. Additionally, the resolution indicates the amount of noise induced into our measurements due to the probe geometry. A dynamic model can only account for one sample and offer a comparison

to later samples; at this point the correction for tracker bias may be applied and see how closely we can resolve the actual position.

4.1 Modeling the Accuracy and Resolution of Beacons

Each tracker measurement for a beacon's position is a value from an underlying distribution of possible values. Thus, we can associate a random variable X with the beacon position. The normalized distribution of values reported by the tracker is referred to as the probability density function (PDF) of X . Accuracy of the beacon can be related to the mean value of X , while a measure of resolution is the variance X .

An intuitive measure of the mean and variance of X are the sample mean and variances of X , respectively. However, for a finite number of samples, those intuitive measures can provide biased estimates [12]. Unbiased estimates can be obtained from the underlying PDF instead.

While the PDF of X will often resemble a normal distribution, it may be significantly different, requiring a separate description. In this case, a useful way to describe the PDF is the Gram-Charlier expansion. The PDF, $p(x)$, is given by

$$p(x) = p_o(x) \left[1 + \sum_{n=1}^N \frac{c_n}{n!} H_n(x) \right], \quad (2)$$

where $p_o(x)$ is the closest normal distribution to $p(x)$, obtained by least-squares fitting [12]. Coefficients c_n measure the departure from the closest normal distribution and are to be determined. The $H_n(x)$ are Hermite polynomials [13]. After fitting the PDF data of one beacon to this model, $p(x)$ is completely determined and can be used in the computation of the mean and variance of X .

If we denote the mean, or first moment of the PDF of X , as the expected value of X , or $E[X]$, $E[X]$ can be estimated from the PDF as

$$E[X] = \int_{-\infty}^{\infty} x p(x) dx, \quad (3)$$

The accuracy is then defined as

$$Accuracy = | E[X] - nominal\ position |, \quad (4)$$

where the nominal position is the value of the known location, confirmed independently. Similarly, if we denote the variance of the random variable X by S_x^2 , it is defined as

$$S_x^2 = E[X^2] - E[X]^2, \quad (5)$$

where $E[X^2]$ is the second moment of X given by

$$E[X^2] = \int_{-\infty}^{\infty} x^2 p(x) dx. \quad (6)$$

The resolution of a beacon is defined as the standard deviation of the random variable, σ_x .

4.2 Modeling the Accuracy of a Probe

Because a probe is an assembly of beacons, it can also be modeled as a random variable, Z . This relationship between probes and beacons is expressed as

$$Z = \sum_{i=1}^N Y_i = \sum_{i=1}^N w_i X_i, \quad (7)$$

where N is the number of beacons on the probe, X_i are the random variables associated with each beacon, and w_i characterizes the distance of each beacon from the probe centroid. In the particular case where the beacons are equidistant from the probe centroid, the case of symmetrical probes, w_i equals $1/N$.

The accuracy of a probe, defined as the expected value of random variable Z , $E[Z]$, is given by

$$E[Z] = \sum_{i=1}^N w_i E[X_i]. \quad (8)$$

For symmetrical probes, $E[Z]$ is given by

$$E[Z] = \frac{1}{N} \sum_{i=1}^N E[X_i] = E[X_i]. \quad (9)$$

Therefore, the accuracy for a probe should be constrained to the accuracy of a single beacon. Additionally, this accuracy in position can be expressed in an identical manner to equation 4.

Without loss of generality, the frame of reference of the probe is defined to be oriented identically to the tracker reference frame. The origin of the probe reference frame is given by the centroid of the beacons on the probe. Modeling the accuracy in orientation of a probe as a function of the probe geometry, its number of beacons, and the beacon characteristics, is equivalent to finding the relationship between the two reference frames using measurements of the positions of individual beacons. After the probe origin has been translated to the origin of the tracker, the rotation between the two reference frames is best described by quaternions [14][15]. The unit quaternion representing the best rotation of the probe has been shown to be the eigenvector associated with the most positive eigenvalue of a symmetric 4x4 matrix [15]. Horn also shows that the elements of this matrix can be expressed as the combination of sums of products corresponding of corresponding positions of the points [15]. Thus, the unit vector describing the orientation of the probe can be measured.

The measured angle of rotation of the probe can be obtained by taking the dot product of two unit vectors, one before rotation and the other after rotation of the probe. The unit vector angle,

$$E[\theta] = \tan^{-1} \frac{A}{D}. \quad (10)$$

denoted as q , is a random variable because it is a sample for an underlying distribution of possible measured angles. However, the expected value of q , $E[q]$, will be related to the accuracy of a single beacon. This relationship is expressed as
where A is the accuracy of a single beacon, determined as described in equation 4, and D is the distance of the beacon furthest from the probe centroid. The orientation accuracy of a probe can then expressed as

$$\text{Accuracy} = |E[\theta] - \text{nominal angle}|. \quad (11)$$

4.3 Modeling the Resolution of a Probe

Given that the random variables associated with the beacon position are independent and identically distributed (iid), the variance of the random variables is constant. We can therefore relate the variance of the random variable associated with the probe to that of the individual beacons. This probe variance, referred to as s_z^2 , is then given by

$$s_z^2 = \sum_{i=1}^N s_{y_i}^2 = \sum_{i=1}^N w_i^2 s_x^2. \quad (12)$$

For symmetrical probes, s_z^2 is equal to

$$s_z^2 = \frac{1}{N} s_x^2. \quad (13)$$

From equation 12, we see that as the number of beacons on a probe increases, the variance of Z decreases, and thus the resolution of the probe in position increases. By knowing the beacon variance, s_x^2 , the number of beacons N needed to obtain a desired probe resolution, s_z , can be calculated. If N , the maximum number of beacons allowable, is imposed by other constraints such as frame rate requirements and cost, either a tradeoff needs to be made between resolution and frame rate or new tracking schemes may be considered.

The resolution in orientation, Δq , of a probe is determined by the resolution in position of the beacon on the probe which is furthest away from the centroid. If D is the distance of this beacon to the centroid, and s_x is its resolution in position, the resolution in orientation is given by

$$\Delta q = 2 \tan^{-1} \frac{s_x}{D}. \quad (14)$$

This method of modeling resolution in orientation is an intuitive method, giving an approximation of the achievable resolution. A rigorous derivation based on the position of all markers on the probe is currently under investigation.

5. Experimental Methodology

For the measurements presented, translation and rotation stages with accuracies of 0.01mm and 0.0625 deg respectively were used. Each probe was fastened to the assembly of stages, which was secured on an optical table. All displacements were made from the origin of the stage referred to as (x,y,z) equal to $(0,0,0)$. All measurements were made in the center of the field of view of the tracker at approximately 4.5m from the sensing cameras. The minimum number of data samples required for our measurements to have at least 99% probability of correctness for

the inference that the error on the largest estimated accuracy is 0.02 mm, was calculated to be 374.

Theoretical values for accuracy and resolution were based upon results obtained from one beacon, as explained in section 4. Our theoretical values for accuracy and resolution in position are the experimental values obtained with a single beacon. For accuracy and resolution in orientation, theoretical values are calculated based upon the beacon position farthest from the probe centroid, as detailed in equations 10, 11 and 14.

The probes that were used in the experiment were a one beacon probe, a three beacon probe, a six beacon probe, and a six beacon probe from Northern Digital (referred to as six beacon off-shelf). Illustrations of the probes are found at the end of this document.

5.1. Accuracy in Position

Although precision translation and rotation stages were used, the axis of translation could not be precisely oriented parallel to one of the axes of the OPTOTRAK reference frame without special purpose measurement instruments. Therefore, while we indicate position accuracy in x, y, and z, which refer closely to the tracker reference frame, all measures of accuracy were performed according to Cartesian distance measurements to insure the highest accuracy of the measurements.

Each probe was first positioned at the origin of the experimental setup. Here, 1,000 samples of x, y, and z position data were acquired. This data was averaged for each direction and combined to form a vector indicating the location of the probe as reported by the tracker. This procedure was repeated ten times, totaling 10,000 samples. The ten vectors resulting from the procedure were stored for later usage.

The probe was then moved 10 mm (the nominal distance) along the selected axis of translation. At this final location, the previously described procedure was repeated. Then, the difference between the position vector at the origin and the position vector at the displacement for each of the ten vector pairs was computed. The magnitudes of these vectors were computed, then averaged to give the overall measured distance of translation. The accuracy in position was then computed as the absolute difference between the measured and nominal distance of translation (as described in equation 4).

5.2 Accuracy in Orientation

The OPTOTRAK reports pitch, yaw, and roll of the probe, referred to as β , α , and γ , respectively. These angles are reported with respect to the OPTOTRAK frame of reference. The required unit vector that represents the probe orientation in the OPTOTRAK reference frame (see section 4.2) is computed by applying the pitch, yaw, and roll transformations to a unit vector selected in the OPTOTRAK frame of reference. The transformation matrices are:

$$Pitch = \begin{bmatrix} 1 & 0 & 0 \\ 0 & \cos b & \sin b \\ 0 & \sin b & \cos b \end{bmatrix} Yaw = \begin{bmatrix} \cos a & 0 & \sin a \\ 0 & 1 & 0 \\ -\sin a & 0 & \cos a \end{bmatrix} Roll = \begin{bmatrix} \cos g & -\sin g & 0 \\ \sin g & \cos g & 0 \\ 0 & 0 & 1 \end{bmatrix}. \quad (14)$$

Given these transformations, accuracy in orientation was found in a manner similar to the procedure mentioned earlier for finding accuracy in position. The only difference between the procedures was that the displacement was angular as opposed to linear. The angular displacement was approximately 6 degrees in all three directions, with 1,000 data samples taken at the origin and displacement. The accuracy in orientation was then computed as the absolute difference between the measured and nominal angles of rotation (as described in equation 11).

5.3 Resolution Measurements

The PDF of the random variable associated with a single beacon was estimated using 1,000 samples of position data obtained at the stage origin. The data was then fit to a Gaussian distribution using the sample mean and standard deviation (see Figure 1 below). Furthermore,




Figure 1: Fitted Gaussian for 1,000 data points and Measure of Fit (Residual) Plot

because the PDF of a single beacon was close to a normal distribution, we concluded that the PDF for a probe with N beacons converged to a normal distribution, as specified by the Central Limit Theorem [16]. From the fitted Gaussian PDF, we determined theoretical values for the resolution of all probes in position and orientation.

To measure resolution in position, we followed a methodology similar to that of our accuracy measurements. At the origin, 1,000 samples of x , y , and z data were taken. In this case, a sample variance of the data in each of the directions was computed, then averaged to obtain a measure of variance in each direction. This process was repeated ten times, with each of the ten average variances collected and averaged to form an overall variance measure. Taking the square root of the overall variance gave our resolution measure. The same procedure was applied for determining resolution in orientation, with differences being in the data collected (orientation instead of position).

A complete listing of our data results, along with plots of all experimental data trends may be found at the end of this document.

6. Results and Discussion

In the proposed model of accuracy in position (equation 9), the accuracy of a probe is limited by the accuracy of a single beacon on the probe. Results obtained for the one, three, and six beacon probes closely follow our theoretical predictions. However, the six beacon off-shelf probe deviated from our theoretical predictions in all directions. In the x and z directions, the accuracy of the off-shelf probe was better than that of the probes we constructed. But, in the y direction, the accuracy of the off-shelf probe was worse. A closer look at the results reveals that the accuracy of the off-shelf probe, while not following our predictions, remained essentially the same for all directions while the accuracy of the probes we built fluctuated. We attribute this to the construction of the Northern Digital probe. The LEDs on the probe are permanently attached and recessed into the probe, providing an extremely rigid structure. In addition, this configuration allows the six beacons to remain as close to co-planar as possible. These characteristics add to the overall precision of the probe construction and, therefore, leads to an increasingly precise performance.

The proposed model for resolution in position states that as the number of beacons increases, resolution increases (equation 13). Results obtained with the three beacon and six beacon off-shelf probe followed this prediction. Also, as previously indicated, the results from the off-shelf probe were very precise. The six beacon probe we constructed, however, did not follow the prediction indicated by the model. Again, we think this departure from the model is related to the placement of LEDs on the probe. The three beacon probe behaved according to prediction because its assembly is necessarily planar. A plane can always be found that contains three points. However, for more than three points, a plane that contains all the points may not exist. For this reason, our six beacon probe did not have a completely planar construction and, therefore, did not perform according to its prediction.

Orientation accuracy...Still Working

Orientation resolution...Still Working

7. Conclusions

This report presents a methodology for LED-based optical probe design. Given the requirements of an application for user interactions, accuracy, resolution, tracking speed, and working volume, a probe geometry may be selected and its size and number of beacons calculated to meet specifications. Future work will include the construction of a spherical probe with equally distributed beacons that will satisfy requirements of a specific research application our research laboratory. Tradeoffs discussed here and methodologies presented will guide the design of the probe. Additionally, an effort to construct a rigid body with the LED positions biased to meet the co-planar beacon requirement for planar probes will also be investigated.







Figure 2: Probes used in the experiment (Spherical Probe to be constructed)

Table 1: Summary of accuracy and resolution measurements for position and orientation of one, three, and six beacon probes

| | One Beacon | Three Beacon Probe | | Six Beacon Probe | | Six Beacon Probe (Off Shelf) | |
|-------------------------------|----------------|--------------------|----------------|------------------|----------------|------------------------------|----------------|
| | Measured Value | Predicted Value | Measured Value | Predicted Value | Measured Value | Predicted Value | Measured Value |
| Resolution in X (mm) | | | | | | | |
| Resolution in Y | | | | | | | |
| Resolution in Z | | | | | | | |
| Resolution in Pitch (arc min) | *** | | | | | | |
| Resolution in Yaw | *** | | | | | | |
| Resolution in Roll | *** | | | | | | |
| Accuracy in X (mm) | | | | | | | |
| Accuracy in Y | | | | | | | |
| Accuracy in Z | | | | | | | |
| Accuracy in Pitch (arc min) | *** | | | | | | |
| Accuracy in Yaw | *** | | | | | | |
| Accuracy in Roll | *** | | | | | | |

Figure 4: Experimental Plots (four pages)






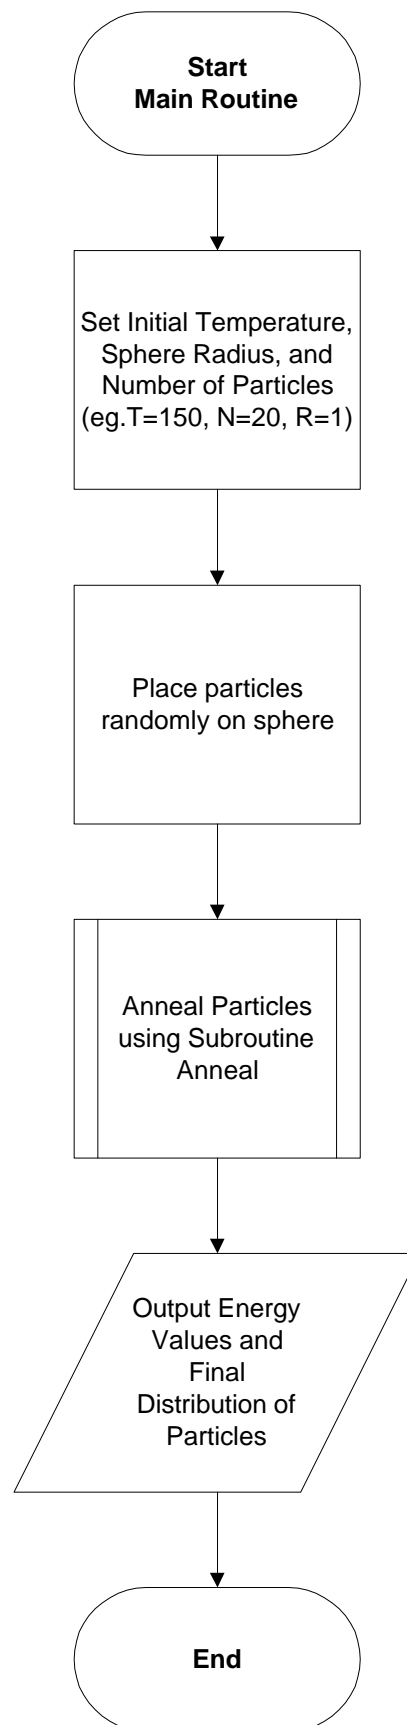
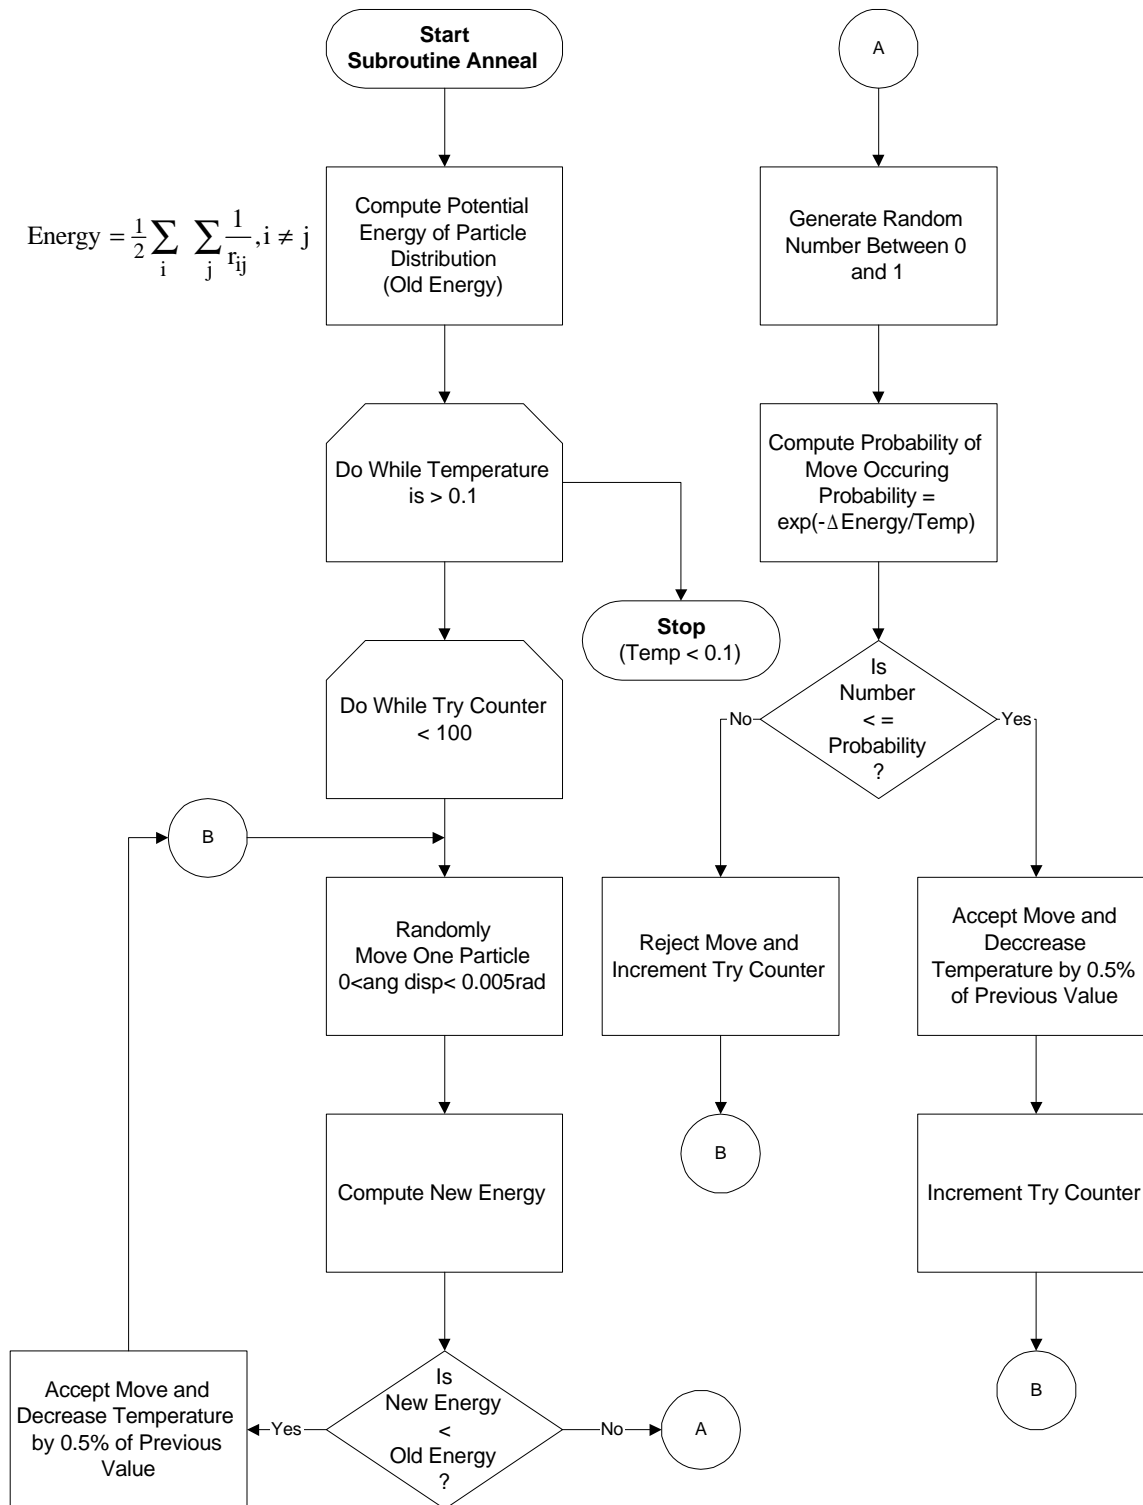


Figure 3: Flowchart of Simulated Annealing Algorithm (two pages)



Anneal Subroutine



References

1. Azuma, R., "Predictive Tracking for Augmented Reality, Ph.D. dissertation, University of North Carolina at Chapel Hill, Department of Computer Science Technical Report TR95-007, 1995.
2. Matzura, A. and S. Seifert, "Virtual Cutting in Medical Data" in Medicine Meets Virtual Reality, Vol. 39, K. S. Morgan, H.M. Hoffman, D. Stredney and S. J. Weghorst, Eds. Amsterdam, Netherlands.: IOS Press, 1997, pp. 420-439.
3. Hinckley, K., R. Pausch, J. Downs, D. Proffitt, and N. Kassell, "The Props-Based Interface for Neurosurgical Visualization" in Medicine Meets Virtual Reality, Vol. 39, K. S. Morgan, H.M. Hoffman, D. Stredney and S. J. Weghorst, Eds. Amsterdam, Netherlands.: IOS Press, 1997, pp. 552-562.
4. Holloway, R., "Registration Errors in Augmented Reality Systems", Ph.D. dissertation, University of North Carolina at Chapel Hill, Department of Computer Science Technical Report TR95-016, 1995
5. Wang, J. "A Real-Time Optical 6-D Tracker for Head Mounted Display Systems", UNC-CS at Chapel Hill, Ph.D. dissertation, TR 90-011.
6. Baillot, Y. and J. Rolland, "Fundamental Principles of Tracking Technology for Virtual Environments", University of Central Florida Technical Report TR96-004, 1996.
7. Welch, G., and G. Bishop, "SCAAT: Incremental Tracking with Incomplete Information," Proceedings of SIGGRAPH (Los Angeles, CA, August 3-8) pp. 333-334, 1997,
8. Wright, D., J. Rolland and A. Kancherla, "Using Virtual Reality to Teach Radiographic Positioning", Radiologic Technology, Vol. 66, no. 4, pp. 167-172, 1995.
9. Rolland, J, D. Wright, and A. Kancherla, "Towards a Novel Augmented-Reality Tool to Visualize Dynamic 3-D Anatomy" in Medicine Meets Virtual Reality, Vol. 39, K. S. Morgan, H.M. Hoffman, D. Stredney and S. J. Weghorst, Eds. Amsterdam, Netherlands.: IOS Press, 1997, pp. 337-348.
10. Northern Digital Incorporated, 1994
11. Press, W, S. Teukolsky, W. Vetterling, and B. Flannery, Numerical Recipes in C: The Art of Scientific Computing, Cambridge University Press, Cambridge, 2nd Ed., 1992.
12. Frieden, B., Probability, Statistical Optics, and Data Testing: A Problem Solving Approach, Springer-Verlag, Berlin, 2nd Ed., 1991
13. Korn, G. and T. Korn, Mathematical Handbook for Scientists and Engineers, McGraw-Hill, New York, 1961
14. Shoemake, K., "Animating Rotation with Quaternion Curves", Proceedings of SIGGRAPH (San Francisco, CA, July 22-26), Vol. 19, No. 3, pp. 245-254, 1985.
15. Horn, B., "Closed-Form Solution of Absolute Orientation Using Unit Quaternions", Journal of the Optical Society of America, Vol. 4, No. 4, April 1987.
16. Montgomery, D. and G. Runger, Applied Statistics and Probability for Engineers, John Wiley & Sons, Inc, Canada, 1994.

Acknowledgments

The authors thank Dr. Boris Z'eldovich for his discussions on the platonic solids and their relation to nature and Mr. Alexei Goon for his assistance with the distribution algorithm. This research was supported, in part, by funding from the National Institute of Health, the Florida Education Fund (McKnight Doctoral Fellowship), and the UCF Office of Graduate Research.